HP StorageWorks
MPxIO for Sun Solaris application notes

Part number: AA-RW7CD-TE
Third edition: August 2006

Legal and notice information

About this document

This section describes the following topics:

- Application notes information
- Intended audience
- Other documentation

Application notes information

These application notes describe the following:

Intended audience

This document is intended for customers who are using Sun Solaris MPxIO with HP storage systems.

Other documentation

A complete library of related documentation is available at:

http://h18006.www1.hp.com/storage/saninfrastructure.html.

For Solaris 8 & 9, the following documentation is available at http://www.sun.com/products-n-solutions/hardware/docs/Network_Storage_Solutions/SAN/san_software/index.htm:

- Sun StorEdge SAN Foundation Software 4.4 Installation Guide
- Sun StorEdge SAN Foundation Software 4.4 Configuration Guide
- Sun StorEdge SAN Foundation Software 4.4.x Release Notes, including Sun StorEdge Traffic Manager for the Solaris Operating System
- Sun StorEdge Traffic Manager Installation and Configuration Guide For the Solaris Operating System
- Sun StorEdge SAN Foundation Software 4.4

For Solaris 10, the Solaris Fibre Channel and Storage Multipathing Administration Guide is available at http://docs.sun.com/app/docs/doc/819-0139.

MPxIO overview

MPxIO provides basic failover and load-balancing capability to HP StorageWorks EVA3000/5000 v4, EVA4000/6000/8000, and the XP line of storage systems. Different storage devices can connect to the same system. Failover to other storage devices is allowed.

New features

This release provides support for Sun Solaris MPxIO multipathing on EVA3000/5000 4.004 and EVA4000/6000/8000 disk arrays.

NOTE:

MPxIO is not supported on versions of the EVA3000/5000 prior to 4.004.

Supported hardware and software

Table 1 describes supported hardware and software.

Table 1 Supported hardware and software

Component	Version
Operating system	Solaris 8, 9, 10
Sun native adapters	SG-XPCI1FC-QF2 (X6767A) SG-XPCI2FC-QF2 (X6768A) SG-XPCI2FC-QF2–Z SG-XPCI1FC-QF4 SG-XPCI2FC-QF4 SG-XPCIE1FC-QF4 (Solaris 10 only) SG-XPCIE2FC-QF4 (Solaris 10 only) SG-XPCI1FC-QL2 SG-XPCI1FC-EM2 SG-XPCI2FC-EM2 SG-XPCIE1FC-EM4 (Solaris 10 only) SG-XPCIE2FC-EM4 (Solaris 10 only)
QLogic adapters	QLA2310F QLA2340 QLA2342 QLA2460 QLA2462
Emulex adapters	LP10000 LP10000DC LP11000 LP11002
Storage Systems	HP StorageWorks EVA3000/5000 4.004 HP StorageWorks EVA4000/6000/8000 HP StorageWorks XP48/512 HP StorageWorks XP128/1024 HP StorageWorks XP10000/12000

Supported multipathing combinations

The combination of MPxIO and Secure Path on the same server can only be supported if FCA2257P is used for Secure Path and Sun native or Emulex adapters are used for MPxIO.

To view more multipathing combinations supported by MPxIO, go to http://www.hp.com/go/sandesignguide. From the web page, select the complete *HP StorageWorks SAN design reference guide* (1-5) and check the table *Sun Solaris multipathing coexistence support.*

Preparing the host system

To set up the host system:

1. For Solaris 8 & 9 and Fibre Channel Adapters (FCAs) from Sun or QLogic:

   If the OS has been installed before adding QLogic based FCAs, you may need to install SUNWqlc and SUNWqlcx first. Go to http://sunsolve.sun.com (and select *Document id 74507*) for further details.

2. Install the necessary software components for the different Solaris versions:

   a. For Solaris 8 and 9, download the Sun SAN Foundation Software (SFS) 4.4.8 or later, SAN_4.4.x_install_it.tar.Z which includes all the patches. Install it with the included `install_it` script, available at http://www.sun.com/storagetek/storage_networking/ by selecting *SAN 4.4 release Software/Firmware Upgrades & Documentation*. To verify which version of SFS has been installed, go to http://sunsolve.sun.com (see *Document id 77230*).

   b. For Solaris 10, the required patches depend on the FCA.
      - With QLogic based FCAs (from Sun or QLogic) SUNWqlc with Patch 119130–16 or later is required
      - With Emulex based FCAs (from Sun or Emulex), SUNWemlxs and SUNWemlxu with Patch 120222–08 or later is required
      - Use `/usr/bin/updatemanager` (see http://www.sun.com/service/sunupdate/) to update to the current Sun SAN patch level

3. Reboot the Host after installing the (software and) patches.

4. Upgrade each FCA to the latest FCA Fcode/OpenBoot if required (typically requires single user mode and no I/O activity):
   - For FCAs from Sun, go to http://sunsolve.sun.com/ or check with Sun if the FCA requires an fcode update and follow the instructions in the patch description in case a patch is available
   - For native Emulex FCAs, check the Emulex web site (http://www.emulex.com) for the latest OpenBoot version for the given FCA. Download the fcode and FCA utilities from that web site and use `/opt/EMLXemlxu/bin/emlxadm` to check and update the fcode
   - For native QLogic FCAs, check the QLogic web site at http://www.qlogic.com for the latest qlc Fcode. Download the fcode and SANSurfer CLI from that web site and use `/opt/QLogic_Corporation/SANsurferCLI/scli` to check and update the fcode

5. Reboot the Host after updating the Fcode on the FCAs

Configure and verify the host

Verifying configuration settings

Verify that all paths to the storage system are present with `cfgadm -al` command. To display the LUNs configured to each path, execute `cfgadm -al -o show_FCP_dev` as shown in the example below.

The following examples are with 4 paths (ports) visible from the array on the host and with 2 LUNs configured on the array for the host. You should normally see 4 paths with EVA3000/ EVA5000/ EVA4000/ EVA6000, however you might also see up to 8 paths with EVA8000.

1. Enter the `cfgadm` command to verify the paths to the LUNs (note that you need to have configured LUNs on the array for the host in order to see the "disk" lines):

```
Example: # cfgadm -al -o show_FCP_dev
Ap_Id                    Type         Receptacle    Occupant        Condition
c2                       fc-fabric    connected     unconfigured    unknown
c2::50001fe1002709f8,0   array-ctrl   connected     unconfigured    unknown
c2::50001fe1002709f8,1   disk         connected     unconfigured    unknown
c2::50001fe1002709f8,2   disk         connected     unconfigured    unknown
c2::50001fe1002709fc,0   array-ctrl   connected     unconfigured    unknown
c2::50001fe1002709fc,1   disk         connected     unconfigured    unknown
c2::50001fe1002709fc,2   disk         connected     unconfigured    unknown
c3                       c-fabric     connected     unconfigured    unknown
c3::50001fe1002709f9,0   array-ctrl   connected     unconfigured    unknown
c3::50001fe1002709f9,1   disk         connected     unconfigured    unknown
c3::50001fe1002709f9,2   disk         connected     unconfigured    unknown
c3::50001fe1002709fd,0   array-ctrl   connected     unconfigured    unknown
c3::50001fe1002709fd,1   disk         connected     unconfigured    unknown
c3::50001fe1002709fd,2   disk         connected     unconfigured    unknown
```

With Solaris 8 or 9, the Occupant status will be displayed as above in a SAN environment. With Solaris 10, you will see a "configured" status.

2. For Solaris 8 and 9, use `cfgadm -c configure c#` command to configure the host to access the storage system as shown in the example below. The value of # is gathered from the output of the `cfgadm -al` previously executed.

   Example:

```
# cfgadm -c configure c2
# cfgadm -c configure c3
```

3. Execute the `cfgadm` command to view configured devices. After the configuration, the following command will show the array controllers as configured:

```
# cfgadm -al
Ap_Id                    Type         Receptacle    Occupant        Condition
c0                       scsi-bus     connected     configured      unknown
c0::dsk/c0t0d0           disk         connected     configured      unknown
c0::dsk/c0t6d0           CD-ROM       connected     configured      unknown
c1                       scsi-bus     connected     unconfigured    unknown
c2                       fc-fabric    connected     configured      unknown
c2::50001fe1002709f8     array-ctrl   connected     configured      unknown
c2::50001fe1002709fc     array-ctrl   connected     configured      unknown
c3                       fc-fabric    connected     configured      unknown
c3::50001fe1002709f9     array-ctrl   connected     configured      unknown
c3::50001fe1002709fd     array-ctrl   connected     configured      unknown
```

NOTE:

Using the option `–o show_FCP_dev` could result in the output of the `cfgadm` command to display the Type `array-ctrl` as unconfigured. This is not a problem. (See example below).

```
# cfgadm -al -o show_FCP_dev
Ap_Id                    Type         Receptacle    Occupant        Condition
c2                       fc-fabric    connected     configured      unknown
c2::50001fe1002709f8,0   array-ctrl   connected     unconfigured    unknown
c2::50001fe1002709f8,1   disk         connected     configured      unknown
c2::50001fe1002709f8,2   disk         connected     configured      unknown
...
```

4. Verify that the LUNs/disks are visible with the `format` command through all the EVA paths.

```
# echo | format
Searching for disks...done
c2t50001FE1002709F8d1: configured with capacity of 1008.00MB
c2t50001FE1002709F8d2: configured with capacity of 1008.00MB
c2t50001FE1002709FCd1: configured with capacity of 1008.00MB
c2t50001FE1002709FCd2: configured with capacity of 1008.00MB
c3t50001FE1002709F9d1: configured with capacity of 1008.00MB
c3t50001FE1002709F9d2: configured with capacity of 1008.00MB
c3t50001FE1002709FDd1: configured with capacity of 1008.00MB
c3t50001FE1002709FDd2: configured with capacity of 1008.00MB

AVAILABLE DISK SELECTIONS:
0. c0t0d0 /pci@1f,4000/scsi@3/sd@0,0
1. c2t50001FE1002709F8d1 /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709f8,1
2. c2t50001FE1002709F8d2 /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709f8,2
3. c2t50001FE1002709FCd1 /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709fc,1
4. c2t50001FE1002709FCd2 /pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709fc,2
5. c3t50001FE1002709F9d1 /pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709f9,1
6. c3t50001FE1002709F9d2 /pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709f9,2
7. c3t50001FE1002709FDd1 /pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709fd,1
8. c3t50001FE1002709FDd2 /pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709fd,2

Specify disk (enter its number)
```

5. Check the multi-pathing information for all LUNs (first example) and either for a given disk/LUN or for the array WWN or a port WWN. The multipathing information is already available although MPxIO has not yet been enabled:

```
# luxadm probe
No Network Array enclosures found in /dev/es
Found Fibre Channel device(s):

Node WWN:50001fe1002709f0  Device Type: Disk device
  Logical Path:/dev/rdsk/c2t50001FE1002709F8d1s2
  Logical Path:/dev/rdsk/c2t50001FE1002709FCd1s2
  Logical Path:/dev/rdsk/c3t50001FE1002709F9d1s2
  Logical Path:/dev/rdsk/c3t50001FE1002709FDd1s2
Node WWN:50001fe1002709f0  Device Type: Disk device
  Logical Path:/dev/rdsk/c2t50001FE1002709F8d2s2
  Logical Path:/dev/rdsk/c2t50001FE1002709FCd2s2
  Logical Path:/dev/rdsk/c3t50001FE1002709F9d2s2
  Logical Path:/dev/rdsk/c3t50001FE1002709FDd2s2

# luxadm display /dev/rdsk/c2t50001FE1002709F8d1s2
DEVICE PROPERTIES for disk: /dev/rdsk/c2t50001FE1002709F8d1s2
  Vendor: HP
  Product ID: HSV210
  Revision: 5100
  Serial Num: Unavailable
  Unformatted capacity: 1024.000 MBytes
```

```
  Read Cache: Enabled
    Minimum prefetch: 0x0
    Maximum prefetch: 0x0
  Device Type: Disk device
  Path(s):

/dev/rdsk/c2t50001FE1002709F8d1s2
/devices/pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709f8,1:c,raw
  LUN path port WWN:          50001fe1002709f8
  Host controller port WWN:   210000e08b1759bd
  Path status:                O.K.
/dev/rdsk/c2t50001FE1002709FCd1s2
/devices/pci@1f,4000/QLGC,qla@4/fp@0,0/ssd@w50001fe1002709fc,1:c,raw
  LUN path port WWN:          50001fe1002709fc
  Host controller port WWN:   210000e08b1759bd
  Path status:                O.K.
/dev/rdsk/c3t50001FE1002709F9d1s2
/devices/pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709f9,1:c,raw
  LUN path port WWN:          50001fe1002709f9
  Host controller port WWN:   10000000c93d5cee
  Path status:                O.K.
/dev/rdsk/c3t50001FE1002709FDd1s2
/devices/pci@1f,4000/lpfc@5/fp@0,0/ssd@w50001fe1002709fd,1:c,raw
  LUN path port WWN:          50001fe1002709fd
  Host controller port WWN:   10000000c93d5cee
  Path status:                O.K.
```

Errors for Solaris 8, 9, and 10

You will see error messages similar to the following for each EVA controller for LUN 0 on the console and in `/var/adm/messages`. You can ignore these error messages.

```
May 18 15:10:30 simonb fcp:  [ID 305930 kern.warning] WARNING: fp1:  no
driver for device @w50001fe1002709e8,0:

May 18 15:10:30 simonb compatible:  scsiclass,0c.vHP.pHSV210.r5100
scsiclass,0c.vHP.pHSV210 scsa,0c.bfcp scsiclass,0c scsiclass

May 18 15:10:30 simonb scsi:  [ID 243001 kern.warning] WARNING:
/pseudo/fcp@0 (fcp0):

May 18 15:10:30 simonb Create devinfo failed.
```

Enabling MPxIO

To enable MPxIO for HP storage devices, the appropriate information for `device-type-scsi-options-list` need to be added in the `/kernel/drv/scsi_vhci.conf` file.

1. Use a text editor to change the configuration file, e.g.
   ```
   # vi /kernel/drv/scsi_vhci.conf
   ```
2. With Solaris 8 or 9 only you need to change the mpxio-disable parameter to the following:
   ```
   mpxio-disable="no";
   ```
3. Verify that the load-balance and auto-failback parameters are set to the following values (the default setting may vary depending on the patch level / OS version):
   ```
   load-balance="round-robin";
   auto-failback="disable";
   ```
4. Add the following lines to cover all HP arrays:

```
device-type-scsi-options-list =
"HP      HSV101", "symmetric-option",
"COMPAQ  HSV111", "symmetric-option",
"HP      HSV2", "symmetric-option",
"HP      OPEN", "symmetric-option";
symmetric-option = 0x1000000;
```

NOTE:

Enter six spaces after "HP" and two spaces after "COMPAQ".

5. Depending on the Solaris version, you need to do the following:

   a. With Solaris 8 or 9 you need to run a reconfiguration `reboot` in order to activate the changes:
   ```
   # reboot -- -r
   ```
   b. With Solaris 10 you need to run the `stmsboot` command and confirm the reboot:
   ```
   # stmsboot –e
   ```

Verify the MPxIO configuration

After the system has rebooted, use the `format` or `luxadm` command to verify that the MPxIO enablement has been successful.

NOTE:

In the `format` or `luxadm` output, controllers c2 and c3 no longer show up. A single, virtual controller path such as c4 now displays the MPxIO–configured LUNs. The `cfgadm` output does NOT change after the MPxIO enablement.

1. Verify with either `format` or `luxadm` probe that only one device file is shown for each LUN:

```
# echo | format
Searching for disks...done

c4t600508B400102E640000B000081C0000d0: configured with capacity of 1008.00MB
c4t600508B400102E640000B00008100000d0: configured with capacity of 1008.00MB

AVAILABLE DISK SELECTIONS:
0. c0t0d0 /pci@1f,4000/scsi@3/sd@0,0
1. c4t600508B400102E640000B000081C0000d0   /scsi_vhci/ssd@g600508b400102e640000b000081c0000
2. c4t600508B400102E640000B00008100000d0   /scsi_vhci/ssd@g600508b400102e640000b00008100000
Specify disk (enter its number):

# luxadm probe
No Network Array enclosures found in /dev/es

Found Fibre Channel device(s):
Node WWN:50001fe1002709f0  Device Type:Disk device
  Logical Path:/dev/rdsk/c4t600508B400102E640000B000081C0000d0s2
Node WWN:50001fe1002709f0  Device Type:Disk device
  Logical Path:/dev/rdsk/c4t600508B400102E640000B00008100000d0s2
```

2. Check the different paths with the `luxadm display` command (either for a given disk or for the array WWN or a port WWN):

```
# luxadm display 50001fe1002709f0
DEVICE PROPERTIES for disk: 50001fe1002709f0
```

```
Vendor:               HP
Product ID:           HSV210
Revision:             5100
Serial Num:           Unavailable
Unformatted capacity: 1024.000 MBytes
Read Cache:           Enabled
  Minimum prefetch:   0x0
  Maximum prefetch:   0x0
Device Type:          Disk device
Path(s):

/dev/rdsk/c4t600508B400102E640000B000081C0000d0s2
/devices/scsi_vhci/ssd@g600508b400102e640000b000081c0000:c,raw
Controller                  /devices/pci@1f,4000/QLGC,qla@4/fp@0,0
Device Address              50001fe1002709f8,2
Host controller port WWN    210000e08b1759bd
Class                       primary
State                       ONLINE
Controller                  /devices/pci@1f,4000/QLGC,qla@4/fp@0,0
Device Address              50001fe1002709fc,2
Host controller port WWN    210000e08b1759bd
Class                       primary
State                       ONLINE
Controller                  /devices/pci@1f,4000/lpfc@5/fp@0,0
Device Address              50001fe1002709f9,2
Host controller port WWN    10000000c93d5cee
Class                       primary
State                       ONLINE
Controller                  /devices/pci@1f,4000/lpfc@5/fp@0,0
Device Address              50001fe1002709fd,2
Host controller port WWN    10000000c93d5cee
Class                       primary
State                       ONLINE
...
```